

Amendments to the Claims:

Listing of Claims:

1-3. (Canceled)

4. (Previously Presented) A routing method in a data processing system comprising the steps of:

- receiving a data packet;
- retrieving a destination address from the data packet;
- hashing the destination address to determine a table index into a table in a computer readable medium;

- reading a target address from a table entry using the table index, wherein the target address has been related to and stored in the table entry based on a computed value from a relation computation using the table index and the target address as operands in the relation computation;

- modifying the data packet by storing the target address in the data packet; and
- transmitting the modified data packet.

5. (Previously Presented) The method of claim 4 further comprising a step of relating a particular table entry to a target address in which:

- for each target address in the set of target addresses, generating a computed value using the table index for the particular entry and a target address as operands in the relation computation to obtain a set of computed values;

- choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

- determining a related target address for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related-target address as input.

6. (Currently Amended) The method of claim 5 further comprising:

- obtaining a set of target addresses;

for each table entry, relating a target address from the set of target addresses to a table entry such that each table entry is related with only one target address; and

for each table entry, storing in a table entry ~~its related~~ the target address related with the table entry.

7. (Currently Amended) A method in a data processing system for mapping a source identifier to a target identifier in a set of target identifiers, the method comprising the steps of:

a processor managing a data structure in a computer readable medium, wherein each entry ~~within~~ in the data structure stores information associated with a single target identifier and wherein a single target identifier is related to at least one ~~or more~~ entry location in the data structure ~~locations~~;

the processor hashing the source identifier to a location identifier of an entry in the data structure;

the processor retrieving information associated with the target identifier from the entry in the data structure using the location identifier; [[and]]

the processor obtaining a mapped target identifier from the retrieved information associated with the target identifier; and

a source identified by the source identifier using a target identified by the mapped target identifier in a distributed computing environment,

wherein a processing speed with which the source identifier is mapped to the mapped target identifier is independent of a total number of target identifiers in the set of target identifiers.

8. (Previously Presented) The method of claim 7, wherein the method for mapping the source identifier to the target identifier is stable with respect to changes in the set of target identifiers.

9. (Currently Amended) A method in a data processing system for mapping a source identifier to a target identifier, the method comprising the steps of:

a processor hashing the source identifier to determine a table index into a table in a computer readable medium; [[and]]

the processor reading the target identifier from a table entry in the table using the table index, wherein the target identifier has been related to and stored in the table entry based on a computed value from a relation computation using the table index and the target identifier as operands in the relation computation; and

a source identified by the source identifier using a target identified by the target identifier in a distributed computing environment.

10. (Currently Amended) The method of claim 9, wherein the target identifier is in a set of target identifiers, further comprising a step of the processor relating a particular table entry in the table to a target identifier in the set of target identifiers in which:

for each target identifier in the set of target identifiers, generating a computed value using the table index for the particular table entry and a target identifier as operands in the relation computation to obtain a set of computed values;

choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

determining a related target identifier for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related target identifier as input.

11. (Currently Amended) The method of claim 10, further comprising, prior to the step of the processor reading the target identifier from the table entry:

the processor obtaining a set of target identifiers;

for each table entry in the table, the processor relating a target identifier from the set of target identifiers to a table entry such that each table entry is related with only one target identifier; and

for each table entry, the processor storing in a table entry ~~its related~~ the target identifier related with the table entry.

12. (Currently Amended) The method of claim 10, further comprising:

the processor dynamically removing a target identifier from a set of target identifiers to obtain a modified set of target identifiers;

for each table entry in the table previously related to the removed target identifier, the processor newly relating a target identifier from the modified set of target identifiers to a table entry in the table such that each table entry is related with only one target identifier; and

for each table entry in the table previously related to the removed target identifier, the processor storing in a table entry ~~its newly related~~ the target identifier newly related with the table entry.

13. (Currently Amended) The method of claim 10, further comprising:

the processor dynamically adding a target identifier to a set of target identifiers to obtain a modified set of target identifiers;

for each table entry in the table, the processor relating a target identifier from the modified set of target identifiers to a table entry such that each table entry is related with only one target identifier; and

for each table entry in the table, the processor storing in a table entry ~~its related~~ the target identifier related with the target entry if ~~its related~~ the target identifier related with the target entry differs from a target identifier previously stored in the table entry.

14. (Currently Amended) The method of claim 9, wherein the relation computation further comprises:

the processor receiving the table index and the target identifier as operands for the relation computation;

the processor hashing the table index to generate a first hash value;

the processor hashing the target identifier to generate a second hash value; and

the processor hashing the first hash value and the second hash value to generate $[[a]]$ the computed value.

15. (Currently Amended) The method of claim 9, further comprising:

the processor obtaining a set of target identifiers, wherein each target identifier in the set of target identifiers identifies a computational resource such that each target identifier is related with only one computational resource.

16. (Currently Amended) The method of claim 9, further comprising:
the processor associating a computational resource with a subset of a set of target identifiers, wherein each target identifier in the set of target identifiers is related with only one computational resource, wherein each target identifier in the subset of target identifiers identifies the computational resource, and wherein a size of the subset of target identifiers is proportional to a computational capacity of the computational resource.
17. (Previously Presented) The method of claim 9, wherein the source identifier is a network protocol address.
18. (Previously Presented) The method of claim 9, wherein the target identifier is a network physical address.
19. (Previously Presented) The method of claim 9, wherein the target identifier is a Uniform Resource Identifier (URI).
- 20-41. (Canceled)
42. (Currently Amended) A method in a data processing system for mapping a source identifier to a target identifier, the method comprising steps of:
a processor hashing the source identifier to determine a location identifier of an entry in a data structure in a computer readable medium; [[and]]
the processor reading information associated with the target identifier from the entry in the data structure using the location identifier, wherein the information associated with the target identifier has been related to and stored in the entry based on a computed value from a relation computation using the location identifier and the target identifier as operands in the relation computation; and
a source identified by the source identifier using a target identified by the target identifier in a distributed computing environment.

43. (Currently Amended) The method of claim 42, wherein the target identifier is in a set of target identifiers, further comprising a step of the processor relating a particular entry in the data structure to a target identifier in the set of target identifiers in which:

for each target identifier in the set of target identifiers, generating a computed value using the location identifier of the particular entry in the data structure and a target identifier as operands in the relation computation to obtain a set of computed values;

choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

determining a related target identifier for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related target identifier as input.

44. (Currently Amended) The method of claim 43, further comprising, prior to the step of reading the information associated with the target identifier from the entry in the data structure:

the processor obtaining a set of target identifiers;

for each entry in the data structure, the processor relating a target identifier from the set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

for each entry in the data structure, the processor storing in an entry information associated with ~~its related~~ the target identifier related with the entry.

45. (Currently Amended) The method of claim 43, further comprising:

the processor dynamically removing a target identifier from a set of target identifiers to obtain a modified set of target identifiers;

for each entry in the data structure previously related to the removed target identifier, the processor newly relating a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

for each entry in the data structure previously related to the removed target identifier, the processor storing in an entry information associated with ~~its newly related~~ the target identifier newly related with the entry.

46. (Currently Amended) The method of claim 43, further comprising:
the processor dynamically adding a target identifier to a set of target identifiers to obtain a modified set of target identifiers;

for each entry in the data structure, the processor relating a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

for each entry in the data structure, the processor storing in an entry information associated with ~~its related~~ the target identifier related with the entry if ~~its related~~ the target identifier related with the entry differs from a target identifier previously related to the entry in the data structure.

47. (Currently Amended) The method of claim 42, wherein the relation computation further comprises:

the processor receiving the location identifier and the target identifier as operands for the relation computation;

the processor hashing the location identifier to generate a first hash value;

the processor hashing the target identifier to generate a second hash value; and

the processor hashing the first hash value and the second hash value to generate [[a]] the computed value.

48. (Currently Amended) The method of claim 42, further comprising:

the processor obtaining a set of target identifiers, wherein each target identifier in the set of target identifiers identifies a computational resource such that each target identifier is related with only one computational resource.

49. (Currently Amended) The method of claim 42, further comprising:

the processor associating a computational resource with a subset of a set of target identifiers, wherein each target identifier in the set of target identifiers is related with only one computational resource, wherein each target identifier in the subset of target identifiers identifies the computational resource, and wherein a size of the subset of target identifiers is proportional to a computational capacity of the computational resource.

50. (Currently Amended) The method of claim 42, further comprising:
 the processor retrieving the target identifier using the information associated with the target identifier; and
 the processor performing a computational process on a computational resource identified by the target identifier.
51. (Previously Presented) The method of claim 50, wherein the computational resource identified by the target identifier is a memory resource.
52. (Previously Presented) The method of claim 50, wherein the computational resource identified by the target identifier is a data processing system.
53. (Previously Presented) The method of claim 42, wherein the information associated with the target identifier comprises the target identifier.
54. (Previously Presented) The method of claim 42, wherein the data structure is a table, and the location identifier is a table index.
55. (Previously Presented) The method of claim 42, wherein the source identifier is a network protocol address.
56. (Previously Presented) The method of claim 42, wherein the target identifier is a network physical address.
57. (Previously Presented) The method of claim 42, wherein the target identifier is a Uniform Resource Identifier (URI).
58. (Currently Amended) A data processing system that enables a mapping of a source identifier to a target identifier, the data processing system comprising:
 a processor; and
 a memory comprising instructions, which when executed by the processor, perform actions comprising:

hashing the source identifier to determine a location identifier of an entry in a data structure in a computer readable medium; [[and]]

reading information associated with the target identifier from the entry in the data structure using the location identifier, wherein the information associated with the target identifier has been related to and stored in the entry based on a computed value from a relation computation using the location identifier and the target identifier as operands in the relation computation; and

a source identified by the source identifier using a target identified by the mapped target identifier in a distributed computing environment.

59. (Currently Amended) The data processing system of claim 58, wherein the target identifier is in a set of target identifiers, further comprising first relating means for relating a particular entry in the data structure to a target identifier in the set of target identifiers wherein:

generating, for each target identifier in the set of target identifiers, a computed value using the location identifier of the particular entry in the data structure and a target identifier as operands in the relation computation to obtain a set of computed values;

choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

determining a related target identifier for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related target identifier as input.

60. (Currently Amended) The data processing system of claim 59, further comprising, prior to the step of reading the information associated with the target identifier from the entry in the data structure:

obtaining a set of target identifiers;

relating, for each entry in the data structure, a target identifier from the set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

storing, in each entry in the data structure, information associated with ~~its related~~ the target identifier related with the table entry.

61. (Currently Amended) The data processing system of claim 59, further comprising:
dynamically removing a target identifier from a set of target identifiers to obtain a modified set of target identifiers;

newly relating, for each entry in the data structure previously related to the removed target identifier, a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

storing, in each entry in the data structure previously related to the removed target identifier, information associated with ~~its newly related~~ the target identifier newly related with the table entry.

62. (Currently Amended) The data processing system of claim 59, further comprising:
dynamically adding a target identifier to a set of target identifiers to obtain a modified set of target identifiers;

relating, for each entry in the data structure, a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and

storing, in each entry in the data structure, information associated with ~~its related~~ the target identifier related with the entry if ~~its related~~ the target identifier related with the entry differs from a target identifier previously related to the entry in the data structure.

63. (Currently Amended) The data processing system of claim 58, wherein the relation computation further comprises:

receiving the location identifier and the target identifier as operands for the relation computation;

hashing the location identifier to generate a first hash value;

hashing the target identifier to generate a second hash value; and

hashing the first hash value and the second hash value to generate $[[a]]$ the computed value.

64. (Currently Amended) The data processing system of claim 58, further comprising:
obtaining a set of target identifiers, wherein each target identifier in the set of target identifiers identifies a computational resource such that each target identifier is related with only one computational resource.
65. (Previously Presented) The data processing system of claim 58, further comprising:
associating a computational resource with a subset of a set of target identifiers, wherein each target identifier in the set of target identifiers is related with only one computational resource, wherein each target identifier in the subset of target identifiers identifies the computational resource, and wherein a size of the subset of target identifiers is proportional to a computational capacity of the computational resource.
66. (Previously Presented) The data processing system of claim 58, further comprising:
retrieving the target identifier using the information associated with the target identifier;
and
performing a computational process on a computational resource identified by the target identifier.
67. (Previously Presented) The data processing system of claim 64, wherein the computational resource identified by the target identifier is a memory resource.
68. (Previously Presented) The data processing system of claim 64, wherein the computational resource identified by the target identifier is a router.
69. (Previously Presented) The data processing system of claim 58, wherein the information associated with the target identifier comprises the target identifier.
70. (Previously Presented) The data processing system of claim 58, wherein the data structure is a table, and the location identifier is a table index.
71. (Previously Presented) The data processing system of claim 58, wherein the source identifier is a network protocol address.

72. (Previously Presented) The data processing system of claim 58, wherein the target identifier is a network physical address.

73. (Previously Presented) The data processing system of claim 58, wherein the target identifier is a Uniform Resource Identifier (URI).

74. (Currently Amended) A computer program product comprising:

a non-transitory computer readable medium including instructions, which when executed in a data processing system, map a source identifier to a target identifier, the instructions including computer program product comprising:

instructions for hashing the source identifier to determine a location identifier of an entry in a data structure in a computer readable medium; [[and]]

instructions for reading information associated with the target identifier from the entry in the data structure using the location identifier, wherein the information associated with the target identifier has been related to and stored in the entry based on a computed value from a relation computation using the location identifier and the target identifier as operands in the relation computation; and

instructions for a source identified by the source identifier to use a target identified by the target identifier in a distributed computing environment.

75. (Currently Amended) The computer program product of claim 74, wherein the target identifier is in a set of target identifiers, further comprising instructions relating a particular entry in the data structure to a target identifier in the set of target identifiers including:

instructions for generating, for each target identifier in the set of target identifiers, a computed value using the location identifier of the particular entry in the data structure and a target identifier as operands in the relation computation to obtain a set of computed values;

instructions for choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

instructions for determining a related target identifier for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related target identifier as input.

76. (Currently Amended) The computer program product of claim 75, further comprising:
instructions for obtaining a set of target identifiers;
instructions for relating, for each entry in the data structure, a target identifier from the set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and
instructions for storing, in each entry in the data structure, information associated with ~~its~~ related the target identifier related with the entry.

77. (Currently Amended) The computer program product of claim 75, further comprising:
instructions for dynamically removing a target identifier from a set of target identifiers to obtain a modified set of target identifiers;
instructions for newly relating, for each entry in the data structure previously related to the removed target identifier, a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and
instructions for storing, in each entry in the data structure previously related to the removed target identifier, information associated with ~~its newly-related~~ the target identifier newly related with the entry.

78. (Currently Amended) The computer program product of claim 75, further comprising:
instructions for dynamically adding a target identifier to a set of target identifiers to obtain a modified set of target identifiers;
instructions for relating, for each entry in the data structure, a target identifier from the modified set of target identifiers to an entry in the data structure such that each entry in the data structure is related with only one target identifier; and
instructions for storing, in each entry in the data structure, information associated with ~~its~~ related the target identifier related with the entry if ~~its-related~~ the target identifier related with the entry differs from a target identifier previously related to the entry in the data structure.

79. (Currently Amended) The computer program product of claim 74, wherein the relation computation further comprises:

instructions for receiving the location identifier and the target identifier as operands for the relation computation;

instructions for hashing the location identifier to generate a first hash value;

instructions for hashing the target identifier to generate a second hash value; and

instructions for hashing the first hash value and the second hash value to generate [[a]] the computed value.

80. (Currently Amended) The computer program product of claim 74, further comprising:
instructions for obtaining a set of target identifiers, wherein each target identifier in the set of target identifiers identifies a computational resource such that each target identifier is related with only one computational resource.

81. (Previously Presented) The computer program product of claim 74, further comprising:
instructions for associating a computational resource with a subset of a set of target identifiers, wherein each target identifier in the set of target identifiers is related with only one computational resource, wherein each target identifier in the subset of target identifiers identifies the computational resource, and wherein a size of the subset of target identifiers is proportional to a computational capacity of the computational resource.

82. (Previously Presented) The computer program product of claim 74, further comprising:
instructions for retrieving the target identifier using the information associated with the target identifier; and
instructions for performing a computational process on a computational resource identified by the target identifier.

83. (Previously Presented) The computer program product of claim 80, wherein the computational resource identified by the target identifier is a memory resource.

84. (Previously Presented) The computer program product of claim 80, wherein the computational resource identified by the target identifier is a data processing system.

85. (Previously Presented) The computer program product of claim 74, wherein the information associated with the target identifier comprises the target identifier.

86. (Previously Presented) The computer program product of claim 74, wherein the data structure is a table, and the location identifier is a table index.

87-89. (Canceled)

90. (Previously Presented) A routing method in a data processing system comprising the steps of:

- receiving a data packet;
- retrieving a destination address from the data packet;
- hashing the destination address to determine a table index into a table in a computer readable medium;
- reading a target address from a table entry using the table index, wherein the target address has been related to and stored in the table entry based on a computed value from a relation computation using the table index and the target address as operands in the relation computation, wherein the relation computation is a nearness function;
- modifying the data packet by storing the target address in the data packet; and
- transmitting the modified data packet.

91. (Previously Presented) The method of claim 90, further comprising a step of relating a particular table entry to a target address in which:

- for each target address in the set of target addresses, generating a computed value using the table index for the particular entry and a target address as operands in the relation computation to obtain a set of computed values;

- choosing a computed value from the set of computed values based upon a mathematical relationship among the set of computed values; and

- determining a related target address for the particular entry based on the chosen computed value, wherein the chosen computed value was computed using the related-target address as input.

92. (Currently Amended) The method of claim 91, further comprising:
- obtaining a set of target addresses;
 - for each table entry, relating a target address from the set of target addresses to a table entry such that each table entry is related with only one target address; and
 - for each table entry, storing in a table entry ~~its related~~ the target address related to the table entry.